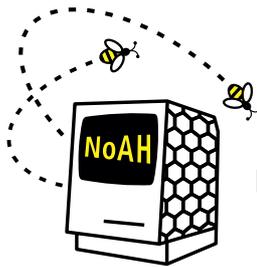


SIXTH FRAMEWORK PROGRAMME
Structuring the European Research Area Specific Programme

RESEARCH INFRASTRUCTURES ACTION



European Network of Affined Honeypots

Contract No. RIDS-011923

D3.1b: Client-side honeypots¹

Abstract:

This document explains how the NoAH infrastructure can be extended with client-side honeypots. Within the project context, we have worked on two different types of honeypots. The first, known as *Shelia*, is fairly traditional and involves a dedicated machine that actively tries to get infected by attacks on client applications. It scans a mailbox and follows every URL contained in the messages. In addition, it opens every attachment with the appropriate application. The *Shelia* client-side honeypot is still novel in the way it detects and contains attack. The second type of client-side honeypot is unlike any other type of honeypot currently available. It runs on the active client machines rather than dedicated honeypots and uses detection techniques that are equivalent to (and thus as reliable as) those of Argos. To cope with the slow-down induced by Argos' taint analysis techniques, the clients' machines only switch to honeypot mode on-demand, or when idle.'

¹This document is based on the paper "Eudaemon: Involuntary and On-Demand Emulation Against Zero-Day Exploits", by H. Bos and G. Portokalidis, also published in the proceedings of EuroSys'08, April 1-4, 2008, Glasgow, Scotland, UK., and also on "Shelia: a client honeypot for client-side attack detection", by Joan Robert Rocaspana, Vrije Universiteit Amsterdam, 2007.

Contractual Date of Delivery	30 September 2008
Actual Date of Delivery	31 October 2008
Deliverable Security Class	Public

The NoAH Consortium consists of:

FORTH	Coordinator	Greece
VU	Principal Contractor	The Netherlands
TERENA	Principal Contractor	The Netherlands
FORTHnet	Principal Contractor	Greece
DFN-CERT	Principal Contractor	Germany
ETHZ	Principal Contractor	Switzerland
VTRIP	Principal Contractor	Greece
ALCATEL	Principal Contractor	France

Contents

1	Introduction	7
2	Shelia: a traditional client-side honeypot	11
2.1	Introduction	11
2.2	Brief Description of Related Work	12
2.3	The Victim - A Profile	13
2.4	A Helicopter View of <i>Shelia</i>	16
2.5	Architecture	17
2.5.1	Email Processor	18
2.5.2	Core Client Emulator	19
2.5.3	Process Monitor Engine	19
2.5.4	Attack Detection Engine	20
2.5.5	Containment Strategy	20
2.6	Attack Detection	21
2.7	Evaluation	23
2.7.1	Protection Effectiveness	23
2.7.2	Results	25
2.8	Performance	26
2.8.1	Micro-benchmarks	26
2.8.2	Macro-benchmarks	27
3	Eudaemon: a good spirit	29
3.1	Related Work	33
3.2	Design	36
3.2.1	Process Possession	37
3.2.2	Process Release	39
3.2.3	Emulator Library	40
3.3	Implementation	40
3.3.1	SEAL: A Secure Emulator Library	41
3.3.2	Possession And Release	45
3.4	Evaluation	49
3.4.1	SEAL	49
3.4.2	Eudaemon	51

CONTENTS

4 Conclusions

53

List of Figures

2.1	The email processor component	19
2.2	Call stack layout (Stack diagram taken from [39])	22
2.3	Micro-benchmark test results	27
2.4	Apache Benchmark. Request-handling capacity decrease measured.	28
3.1	Eudaemon Overview	36
3.2	Process Memory Layout	38
3.3	Process possession: phase 1	46
3.4	Process possession: phase 2	46
3.5	Contents of a /proc/[pid]/maps file - note the presence of <code>libseal</code>	48
3.6	Process release: phase 1	48
3.7	Process release: phase 2	48
3.8	Multiple process possession under low- and high-load	51
3.9	SEAL Execution Loop	52

LIST OF FIGURES

Chapter 1

Introduction

So far, the NoAH project has focused on attacks on servers. In such attacks, the initiative is taken by the attackers. For instance, they may send exploit packets to a web-server to induce a buffer overflow that allows a malicious payload to be executed. The result may be control over the victim machine, elevated privileges, etc.

While attacks on servers remain important, these days we are increasingly witnessing client applications falling victims to malicious servers. Such client-side attacks occur when a user application contacts a server with malicious content. The end result may be exactly the same (for instance, a buffer overflow that leads to a machine that is ‘owned’ by the attacker). The only difference is the interaction that generates this result.

The difference is important for detection. A honeypot like Argos is not able, in and of itself, to find malicious servers. Instead, we need a type of honeypot that explicitly aims to handle this sort of interaction. These honeypots are known as client-side honeypots.

This document explains how the NoAH infrastructure can be extended with client-side honeypots. Within the project context, we have worked on two different types of honeypots. The first, known as *Shelia*, is fairly traditional and involves a dedicated machine that actively tries to get infected by attacks on client applications. It scans a mailbox and follows every URL contained in the messages. In addition, it opens every attachment with the appropriate application. The *Shelia* client-side honeypot is still novel in the way it detects and contains attack. The second type of client-side honeypot, known as *Eudaemon*¹, is unlike any other type of honeypot currently available. It runs on the active client machines rather than dedicated honeypots and uses detection techniques that are equivalent to (and thus as powerful as) those of Argos. To cope with the slow-down induced by Argos’ taint

¹*Eudaemon* was developed in collaboration with the Dutch Deworm project, headed by the Vrije Universiteit

analysis techniques, the clients' machines only switch to honeypot mode on-demand, or when idle.'

Shelia is by far the more mature honeypot. Indeed, it has already been used by other parties and blogged about on various blog sites. Moreover, the Wikipedia entry on client-side honeypots features a separate section on *Shelia*. It is tailored to the most commonly used operating system in the world (Microsoft Windows) and, although still a command-line, research-oriented tool, is designed with usability in mind. *Eudaemon*, in contrast, is very much a prototype. We have been able to evaluate it on Linux, but the current code is neither stable nor very user-friendly.

On the other hand, from a research or novelty perspective, the *Eudaemon* honeypot is probably more interesting. By running on the active client machines rather than dedicated honeypots it circumvents most of the disadvantages of other types of honeypots. In brief, these advantages stem from the fact that a honeypot is not the machine you want to protect.

Eudaemon runs on the exact same machine and uses detection techniques that are the same (and thus as reliable) as those of Argos. In short, *Eudaemon* is a technique that aims to blur the borders between protected and unprotected applications, and brings together honeypot technology and end-user intrusion detection and prevention. *Eudaemon* is able to attach to any running process, and redirect execution to a user-space emulator that will dynamically instrument the binary by means of taint analysis. Any attempts to subvert control flow, or to inject malicious code will be detected and averted. When desired *Eudaemon* can reattach itself to the emulated process, and return execution to the native binary. Selective emulation has been investigated before as a means to heal an attacked program or to generate a vaccine after an attack is detected, by applying intensive instrumentation to the vulnerable region of the program. *Eudaemon* can move an application between protected and native mode at will, e.g., when spare cycles are available, when a system policy ordains it, or when it is explicitly requested. The transition is performed transparently and in very little time, thus incurring minimal disturbance to an actively used system. Systems offering constant protection against similar attacks have also been proposed, but require access to source code or explicit operating system support, and often induce significant performance penalties. We believe that *Eudaemon* offers a flexible mechanism to detect a series of attacks in end-user systems with acceptable overhead. Moreover, we require no modification to the running system and/or installation of a hypervisor, with an eye on putting taint analysis within reach of the average user.

The remainder of this document is organised as follows. In Chapter 2 we discuss the design and implementation of *Shelia*. We show that we are able to capture real attacks and that performance and usability are catered to. In Chapter 3 we discuss *Eudaemon* in detail. As the *Eudaemon* honeypot is more research-oriented, the description is also more research-oriented

(indeed the chapter is directly based on the paper that was published in the ACM SIGOPS EUROSYS conference). We will summarise and conclude in Chapter 4.

Chapter 2

Shelia: a traditional client-side honeypot

Shelia is an intrusion detection system for the client side in the form of a high-interaction honeypot. It comes with a client emulator that scans through a mail folder specified on the command line. Typically, this would be the spam folder. In this folder the client emulator is capable of following every url and opening every attachment. Shelia monitors the processes and generates alerts when the process attempts to execute an invalid operation (i.e., execute a call to change the registry, create files, or attempt specific network operations) from a memory area that is not supposed to be executable code.

2.1 Introduction

Client-side attacks present a problem for attackers. They need to persuade users to visit their malicious server or open files with malicious content. Easily the most commonly used medium for contacting victims is email. For this reason we decided to emulate users by way of their interaction with email.

We have designed and implemented a system for checking automatically all website links and attachments received by email. For each website link or file received, a supervised process is launched to open it: all the API calls made by this process to access the registry, the file-system and the network are validated. This is achieved by means of checking whether the memory address where the API call was made from is contained in an executable region or not. If the address is contained in a region which it is supposed to contain data, an alert is generated. Once an alert is generated, all the data exchanged over the network as well as the information stored in the filesystem is logged by the utility.

Our method of detection has an attractive property; it generates very few (if any) false positives. Any code executing risky operations from the data segment is almost guaranteed to be an attack. The downside is that we limit our scope to a specific set of (extremely harmful) attacks. We do not detect attacks that do not try to inject code to execute.

We should also mention that while our methods were specifically targeted at protecting client application, the engine we use to monitor a process can also be used to monitor a process serving requests such as a http server. The reason is that the detection method is a generic way of capturing malware.

Finally, it is good to point out that the detection method is quite fast, certainly when compared to dynamic taint analysis as employed by Argos and *Eudaemon*. On the other hand, dynamic taint analysis is able to guarantee that not a single instruction of the attack is executed. In *Shelia*, as in most other detection methods (including system call monitoring and anomaly detection), we only detect malicious code by the actions that result from executing the code.

We evaluated *Shelia* using known vulnerabilities published within the Metasploit project and the system designed was able to detect and monitor the execution of most common payloads, independently of the kind of exploit used. In order to evaluate the overhead of a monitored process, we measured the performance of a monitored instance of the Apache server and we compared it with the performance of a non-monitored instance. We used the Apache Benchmark as an evaluation utility and the results showed a minimal decrease in the request handling capacity.

The remainder of this chapter is organised as follows. We start with a brief description of related work in Section 2.2. In Section 2.3 we describe the client-side vulnerabilities. Note that these first two sections could be skipped by readers familiar with these sorts of attacks and defenses. A helicopter view of *Shelia* is provided by Section 2.4, while Section 2.5 provides details about the architecture and Section 2.6 about the way it detects attacks. Evaluations of *Shelia*'s effectiveness and performance are found in Section 2.7 and Section 2.8.

2.2 Brief Description of Related Work

Some well-known client-side honeypots are available and described in the literature. The following list is not meant to be exhaustive. Rather, it is an attempt to place our work on *Shelia* in context.

- HoneyClient: a web browser based (IE/FireFox) high interaction client honeypot implemented by Katy Wang sometime in 2004 or 2005. Snapshots of the registry and file system are taken before any interaction is done. After a website is visited, a long system scan occurs to detect

file/registry changes. As a result, it is very slow. It is an open source project.

- Microsoft's HoneyMonkey: a web browser-based (IE) high interaction client honeypot implemented by Microsoft in 2005 . It is not available for downloading. HoneyMonkey detects attacks on clients by monitoring files, registry, and memory. It consists on an array of virtual machines ranging from completely unpatched Microsoft Windows XP machines up through machines which are completely patched. HoneyMonkey initially crawls a website with a vulnerable configuration. If an attack is detected, the website is re-examined by the next machine in the pipeline. When the end-of-the-pipeline machine is reached, and if the attack is still successful, the URL is upgraded to a zero-day exploit.
- HoneyC: HoneyC is a low interaction client honeypot developed at Victoria University of Wellington by Christian Seifert in 2006. HoneyC is a cross-platform open source framework written in Ruby. Malicious servers are detected by analyzing the web server's responses through the usage of Snort signatures.

2.3 The Victim - A Profile

Client-side attacks rely on vulnerabilities found on desktop applications. There are different attack vectors that can be addressed by an intruder to exploit these vulnerabilities. Microsoft Windows, due to its popularity and large number of attacks, is arguably the most interesting platform for security research. In this section, we will discuss different components that can be targeted by a client-side attack in a Windows environment.

Browser Vulnerabilities

Microsoft Internet Explorer is the most popular browser used for web surfing and it is installed by default on each Windows system. If a vulnerability in this application is identified, a malicious website can be designed to compromise all the workstations from the visiting users. Vulnerabilities in ActiveX controls installed by Microsoft or other vendor software are also exploited via Internet Explorer.

Windows Libraries

Windows libraries are modules that contain functions and data that can be used by other modules such as Windows applications. Windows applications typically leverage a large number of these libraries often packaged as

dynamic-link library (DLL) files to carry out their functions. These libraries usually have the file extension DLL.

These libraries are used for many common tasks such as HTML parsing, image format decoding and protocol decoding. Local as well as remotely accessible applications use these libraries. Thus, a critical vulnerability in a library usually impacts a range of applications from Microsoft and third-party vendors that rely on that library. The most critical issues are the ones that lead to remote code execution without any user interaction when a user visits a malicious web page or reads an email that makes use of the vulnerable library.

ActiveX Components

In order to make web pages and sites more interactive, technologies such as JavaScript, Java applets, etc. have been developed. These are programming languages that let web developers write code that is executed by a web browser locally, that is, on the client workstation.

In order to avoid security problems, the designers of these technologies decided to make use of sandboxed environments. Sandboxing is a popular technique for creating confined execution environments, which could be used for running untrusting programs. A sandbox limits or reduces the access level its applications have. It is a container.

As an example of a sandboxed environment, we can consider JavaScript. By means of this isolated environment, JavaScript scripts are permitted access only to data in the current document or closely related documents (generally those from the same site as the current document). No access is granted to the local file system, the memory space of other running programs, or the operating system's networking layer.

ActiveX is Microsoft's term for a particular kind of software based on the Component Object Model (COM). ActiveX controls are highly portable COM objects, used extensively throughout Microsoft Windows platforms and, especially, in web-based applications. COM objects, including ActiveX controls, can invoke each other locally and remotely through interfaces defined by the COM architecture. The COM architecture allows for interoperability among binary software components produced in disparate ways. Libraries containing ActiveX controls usually have the file extension OCX.

ActiveX controls can also be invoked from web pages through the use of a scripting language or directly with an HTML OBJECT tag. If an ActiveX control is not installed locally, it is possible to specify a URL where the control can be obtained. Once obtained, the control installs itself automatically if permitted by the browser.

ActiveX controls can be signed or unsigned. A signed control provides a high degree of verification that the control was produced by the signer and

has not been modified. Signing does not guarantee the benevolence of the component, it only ensures that the control originated from the signer.

ActiveX controls do not run in a “sandbox” of any kind. That means that ActiveX controls are binary code capable of taking any action that the user can take. Thus, a critical vulnerability in an ActiveX component could lead an attacker to compromise a workstation. An example could be designing a website that, when downloaded into the client, by means of JavaScript, invokes a vulnerable ActiveX component.

Microsoft Office

Microsoft Office is the most widely used email and productivity suite world-wide. The applications include Outlook, Word, PowerPoint, Excel, Visio, FrontPage and Access. Vulnerabilities in these products can be exploited via the following attack vectors:

- The attacker sends a malicious Office document in an email message. The attack succeeds if the user opens it. This is usually achieved by means of social engineering.
- The attacker hosts the document on a web server or shared folder, and entices a user to browse the web page or the shared folder. Note that Internet Explorer automatically opens Office documents. Hence, browsing the malicious web page or folder is sufficient for the vulnerability exploitation.

Attack Example

In September 2006, a vulnerability (CVE-2006-4868) was reported in Windows Vector Markup Language graphics implementation. Vector Markup Language (VML) is a XML-based language typically used to draw vector graphics, i.e. VML graphics. The vulnerable component was VML rendering library Vgx.dll when locating an overly long fill parameter inside a rect tag on a Web page.

This stack overflow flaw was easily exploitable by means of a website invoking the vulnerable library.

Listing 2.1: Example of a malicious HTML document

```
<!-- Currently just a DoS EAX is controllable and currently
it crashes when trying to move EBX into the location pointed
to by EAX [Shirkdog] -->

<html xmlns:v="urn:schemas-microsoft-com:vml">
<head>
```

```
<object id="VMLRender"
  classid="CLSID:10072CEC-8CC1-11D1-986E-00A0C955B42E">
</object>

<style> v\:* { behavior: url(#VMLRender); } </style>

</head>
<body>

<v:rect style='width:120pt;height:80pt' fillcolor="red">
<v:fill method="AAA.....AAAA"
  angle="-45" focus="100%" focusposition=".5,.5"
  focussize="0,0" type="gradientRadial" />
</v:rect>

</body>
</html>
```

2.4 A Helicopter View of *Shelia*

As previously mentioned, attackers move away from server-side attacks and fast-spreading worms, in favour of stealthy attacks that target clients. For client attacks to be successful, an attacker needs to get access to the client's machine or the data transmitted.

In this document, we focus on client-side intrusion attacks, i.e. attacks where attackers try to get their code executed on client machines. *Shelia* builds on three main pillars:

1. Client emulation: for client-side intrusion detection to be successful, we need to mimic the behaviour of a client in order to attract an attack.
2. Attack detection: we need to be able to determine when an attack is taking place
3. Attack analysis: when an attack is taking place, we would like to gather as much information about the attack as it is possible. For instance, we might download the attack's payload for further analysis.

The objective of *Shelia* was to develop these three pillars of a fairly mature client-side honeypot. That is, we intend to develop a system which is able to deal with all the different clients-side attacks previously presented. A common situation where all these attacks can take place is when an unwary user is reading his/her email, and he/she opens the attached files and/or clicks on the links he/she sees. Email messages are widely used by attackers to attach files which contain malicious code or to persuade any potential

reader to visit malicious websites. Our system tries to emulate the behaviour of this gullible user interacting with his/her email inbox.

Since we also intend to detect zero-day attacks, we have aimed for a high-interaction system. Existing high interaction honeyclients rely on the assumption that, after a client-side attack succeeds, some kind of malware is installed in the client machine. As a consequence, they try to detect an attack by means of revising the system state after a server has been interacted with. For example, determining if a file has been added to the file system. Our approach assumes that before any malware could be installed, some kind of vulnerability must be exploited. Our goal is not to detect changes in the system, but the execution of some kind of shellcode trying to download and install malware.

Thus:

1. *Shelia* simulates the behaviour of a user that it is reading emails in his/her email inbox. The system must process all the attachments received and open them with the associated application. Moreover, it launches the default browser to visit all the links to websites included in the content of the messages.
2. For each process launched, the system must validate its behaviour to determine whether there is an attack in process or not.
3. In case there is an attack, we intend to collect information about it, such as the payload it is trying to download.

2.5 Architecture

To obtain a first approximation to our system, we are going to follow the general architecture of a client-sideypot as a guideline. The system should be divided into the following components:

- A queuer: responsible for creating a list of servers for the client to visit.
- A client: responsible for making a request to the server identified by the queuer.
- An analysis engine: responsible for determining whether an attack has taken place on the client honeypot

In our case, the behaviour that we try to reproduce is that of a user interacting with his/her email inbox. Rather than trying to detect an attack while a client email application is exchanging information with the email server, we check that all the content correctly received by an email utility

is not going to compromise the workstation later. That is to say, when a malicious attachment is opened using a desktop application or when a link to a malicious web site is visited by the user.

Thus, the queuer in our design is not only going to create a list of servers (the web sites that must be visited), but also a list of files to validate.

As we aim both to develop a system able to deal with all information that may be received by email (text documents, images, movies, presentations, URLs, etc.) and to scan for attacks based on unknown user application vulnerabilities, the logical choice is to develop a high interaction client honeypot. Phrased differently, we use real applications to process the attachments and URLs identified by the queuer component, instead of developing our own client emulators.

Last but not least, analysis engines in high interaction client honeypots tend to be slower than in low interaction client honeypots, because their detection techniques are based on examining the system state after interaction with (potentially malicious) servers. Our objective is to implement an analysis engine able to detect attacks in an efficient way, but maintaining a low percentage of false positives. Instead of validating the state of the system after the execution of a client, we are going to monitor the behaviour of the client during its execution.

All the mentioned points lead to the following system components:

- Client emulator:
 - Email processor: responsible for identifying all attachments and URLs received by email.
 - Core client emulator: responsible for invoking the proper application to deal with each kind of file the email processor identifies.
- Attack detection
 - Process monitor engine: responsible for monitoring the actions executed by the client as well as to apply the containment strategy if required.
 - Attack detection engine: responsible for determining if an action executed by the client is considered illegal.
 - Attack log engine: responsible for gathering as much information from the attack as possible.

2.5.1 Email Processor

The email processor is responsible for processing all the messages of a specific inbox. It must identify all the links to websites received in the content of the messages. To achieve this, it must parse every message content seeking

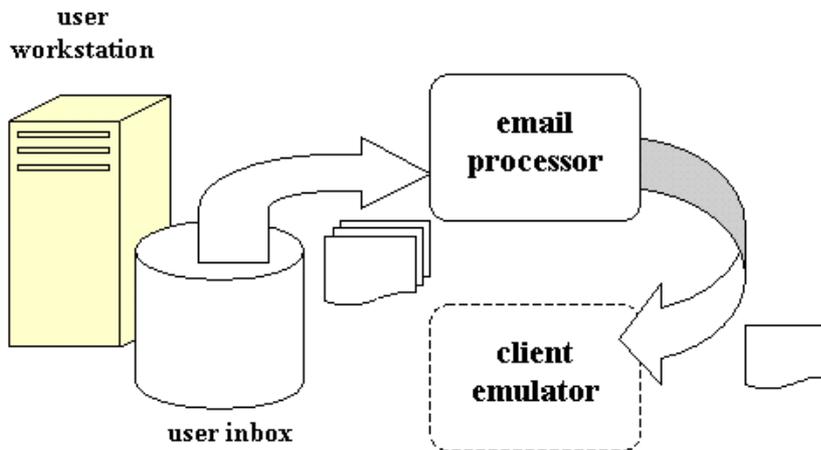


Figure 2.1: The email processor component

for strings that look like URLs. In addition, it must extract all the attachments received in each message. Each URL and attachment extracted from the inbox must be passed to the client emulator. As we are not interested in controlling what happens during the POP3 conversation between a client email utility and a server, we can base our client emulator on an existing application such as Outlook Express, and simply extend it to fulfil our requirements.

2.5.2 Core Client Emulator

The client emulator is invoked by each URL and by any attachment identified by the client emulator. It is responsible for deducing the appropriate application to deal with the information passed from the client emulator and to launch it inside a controlled environment.

When the client emulator is invoked to process an URL, it simply has to open it with the default browser installed on the system. If a file is received as a parameter, before opening it, the associated application must be determined based on the file extension.

2.5.3 Process Monitor Engine

The process monitor engine is responsible for monitoring the actions executed by the applications launched by the client emulator.

Actions that are commonly performed by a successful attack consist of downloading and installing some kind of malware designed to execute some unwanted action. For instance, it may allow a future access to the computer, steal confidential information, serve as a trampoline to attack other hosts,

etc. In order to install malware, modifications in the filesystem must be performed, and, usually, some registry keys must be altered to allow the execution of the malware every time the system starts up. This is the reason why most client honeypots check the state of the system via inspection of changes in the filesystem or in the registry. This is an expensive operation. In contrast, we improve the system performance by only taking into account the actions that access the file system as well as the registry.

Each supervised action must be validated by the detection attack engine. If some action is not validated, the process monitor engine must start logging all the actions performed by the payload. Moreover, it must apply some kind of containment strategy. This aspect will be discussed in more detail later.

2.5.4 Attack Detection Engine

The attack detection engine is responsible for determining whether an action executed by a supervised application is considered illegal or not.

Different techniques can be successfully adapted to our system architecture. A first approximation could be based on the detection of abnormal behaviour. Since the system is going to monitor all the file and registry accesses, we have the possibility to keep information about which files and registry keys are usually accessed by a program. After a learning period, the detection attack engine should be able to determine if an application is following its expected behaviour. The problem with this technique is that it both prone to false positives and difficult to evolve (by continuous training).

Another possibility might be to configure a set of prohibited directories and registry keys. We can assume that a user application doesn't have to save any file into a system directory or that it doesn't have to alter any registry key related to the execution of programs during the system start up. However, this still leaves all other directories vulnerable.

Instead, the criteria we selected for determining whether an action is considered illegal or not is by validating from which memory address the system function was invoked. If this address is not contained into an executable region, an alarm is generated. So, whenever an attacker injects executable code in an application, by whatever means, we will detect an attack as soon as the code tries to perform any of the sensitive operations mentioned above.

2.5.5 Containment Strategy

In case an attack is detected, we would like to gather as much information about it as possible. We expect the shellcode injected to download some kind of malware and install it. We might allow downloading the attack's payload for further analysis, but we probably want to avoid its execution in the normal case. However, sometimes we just want to observe an attack

in the most extensive manner. In that case payload execution should be possible. A *Shelia* command-line option allows administrators to toggle this behaviour to suit their needs.

2.6 Attack Detection

Once an API call is intercepted, we need a mechanism to determine whether this API call must be considered valid or not. The approach that we have followed is based on validating which address the API call was made from. When an application is being executed under the Windows platform, the code sections of the program are loaded in memory regions marked as executable. If a buffer or heap overflow occurs, the payload would have been copied into a writable memory region not supposed to contain instructions.

Our detection attack engine is simply analyzing the stack to determine from which memory address an API call was invoked. Once this address is obtained, it is validated that it is contained in an executable region. If not, an alarm is generated. This calling address is obtained by means of a technique known as stack walking.

Stack Walking

In this subsection we describe the technique of stack walking. The information in this section leans heavily on [39] and [27]. The stack is used for different purposes, but one of the main reasons is to keep track of the point to which each active subroutine should return control when it finishes executing. If, for example, a subroutine *DrawRectangle* calls a subroutine *DrawLine* from four different places, the code of *DrawLine* must have a way of knowing which place to return to.

In addition to the return address, the stack is also used to store the arguments to subroutines, as well as local variables. Information pushed onto the stack as a result of a function call is called a frame. Compilers use a frame base pointer for referencing variables and parameters because their distances from FBP do not change if more push and pop operations are executed.

To understand how stack frames are built, the following C example can be considered:

```
void someFunction(int a, int b, int c, int d) {
    char buffer1[5];
    char buffer2[10];
}

void main() {
    someFunction(a,b,c,d);
}
```

if we disassemble how *someFunction* is invoked, we will get something similar to:

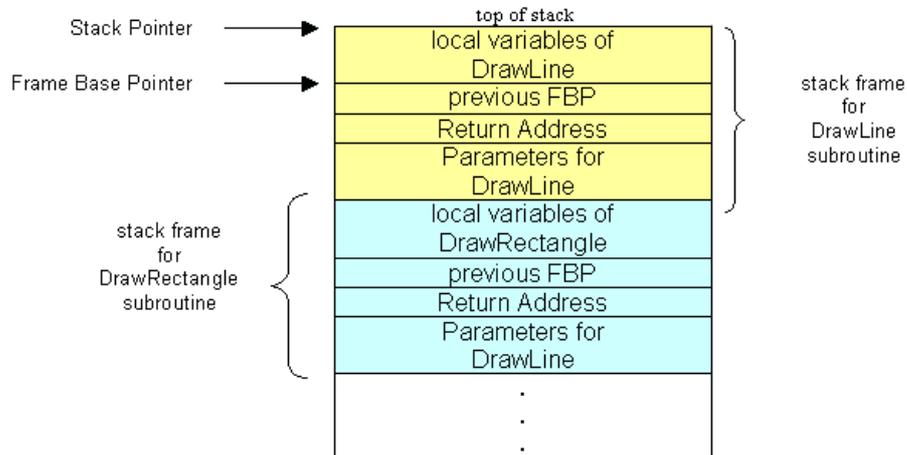


Figure 2.2: Call stack layout (Stack diagram taken from [39])

```

push $d
push $c
push $b
push $a
call someFunction
    
```

First, the parameters to the function are pushed. Secondly, a `CALL` instruction is executed. The address of the next instruction to execute in the calling routine is automatically pushed onto the stack.

Compilers usually generate the following code as the prologue of any function:

```

push ebp
mov  ebp, esp
sub  esp, 20
    
```

The first two lines set up the stack frame. This pushes `EBP` (Extended Base Pointer, Intel’s name for the frame base pointer) onto the stack to save the previous frame base pointer, and it then copies the current `SP` onto `EBP`. This makes `EBP` the new frame base pointer. The third line finds space for local variables by subtracting their size from `ESP`.

When the end of a function is reached, it is executed the epilogue section, which consists on restoring the previous frame base pointer, and returning to the saved return address.

If a subroutine named `DrawLine` is currently running, having just been called by a subroutine `DrawRectangle`, the top part of the call stack might be laid out like figure 2.2.

The technique of stack walking, then, is based on the analysis of the information found in stack frames for determining the calling sequence that execution has followed before arriving to the current routine. Basically, we determine that the context to return to is not contained in the data segment.

2.7 Evaluation

In this section we report our evaluation of the effectiveness and usability of our technique. We conducted experiments to measure both its effectiveness in protecting the system from various attacks, and impact to the system performance.

2.7.1 Protection Effectiveness

We evaluated Shelia using known exploits and payloads available in the Metasploit Framework version 2.7 and 3.1. We first summarise bullet-style all the exploits, operating systems and payloads, while more detailed descriptions follow in the next section. Briefly, then, we evaluated *Shelia* with the following exploits:

- Internet Explorer VML Fill Method Execution (`ie_vml_rectfill`)
- Windows Metafile SetAbortProc Code Execution (`ie_xp_pfv_metafile`)
- Apache Win32 Chunked encoding (`apache_chunked_win32`)

In addition, we tested many other exploits that we will not describe in detail for space reasons. Examples include:

- `windows/browser/ms06_001_wmf_setabortproc`
- `windows/browser/ms06_057_webview_setslice`
- `windows/browser/ms06_013_createtextrange`

The exploits were tested against two vulnerable workstations monitored by Shelia with the following system configurations:

- Windows 2000 SP4
- Windows XP SP2 (DEP protection deactivated)

Each exploits was used to inject the following payloads¹:

- `win32_exec`
- `win32_adduser`
- `win32_bind`
- `win32_downloadexec`

¹The names of the exploits and payloads have changed a little between versions of Metasploit.

2.7.1.1 Vulnerabilities Test Set

We now describe the exploits in a little more detail to give the reader an idea of the comprehensiveness of our evaluation.

VML vulnerability. A stack-based buffer overflow in the Vector Graphics Rendering engine (vgx.dll), as used in Microsoft Outlook and Internet Explorer 6.0 on Windows XP SP2, and possibly in other versions, allows remote attackers to execute arbitrary code via a Vector Markup Language (VML) file with a long fill parameter within a rect tag.

WMF vulnerability. Microsoft Windows WMF graphics rendering engine is affected by a remote code-execution vulnerability. The vulnerability is an inherent defect in the design of WMF files due to the underlying architecture on which they are based.

WMFs are a collection of calls to the Windows GDI (Graphics Device Interface). Code to display a WMF simply passes the calls directly to the GDI. One of the calls allows a user-supplied function to be run in case of print spool cancellation or error. This function is stored in the WMF, and can be anything, for example, a malicious payload. If the function is registered in the GDI by the *setabortproc* GDI call, and an error in rendering the WMF occurs, then the payload will be run.

Apache chunked encoding vulnerability. Apache HTTP Server versions 1.2.2 and later, 1.3 up to and including 1.3.24, and 2.0 up to and including 2.0.36 are vulnerable to a heap buffer overflow in the mechanism that calculates the size of "chunked" encoding. Chunked encoding is a process by which a client generates a variable sized "chunk" of data and notifies the Web server of the data's size before transferring it, so that the Web server can allocate a buffer of the correct size. The Apache HTTP Server has a software flaw that misinterprets the size of incoming data chunks. A remote attacker can use this vulnerability to overflow a buffer and execute arbitrary code or cause a denial of service against the affected Web server.

2.7.1.2 Payload Test Set

Similarly, we describe the payloads to show readers the range of activities that our attacks engage in.

win32_exec. The result of the `win32_exec` payload is the execution of a shell command of the attacker's choosing.

win32_adduser. The result of the `win32_adduser` payload is the adding to the victim host a new user that will give access to the attacker to the compromised workstation.

win32_bind. The result of the `win32_bind` payload is a command shell listening on a port of the attacker's choosing. The payload can be used for remote exploitation after the local exploit has been successfully delivered. If this payload is delivered, an attacker will have access to a command shell running on the victim host.

win32_downloadexec. The `win32_downloadexec` is arguably the most relevant of the four payloads selected for our test. `Win32_downloadexec` downloads a file and executes it on a compromised system. For this test, we chose the notepad application (`notepad.exe`) to be downloaded on the attacked system. It should be noted that the process spawned is executed in the background, so it could only be observed using the Task Manager. Potential malware that may be delivered using `win32_downloadexec` includes spyware, adware, bots, and/or rootkits.

2.7.2 Results

Shelia does not actually prevent exploitation of the system but rather tries to detect the execution of shellcode, by monitoring common API calls used in payloads. As a consequence, the success of the utility does not depend on the kind of vulnerability exploited (buffer overflow, heap overflow, etc.) but on the actions performed by the code injected.

Originally, payloads `win32_exec` and `win32_adduser` couldn't be detected by *Shelia*. The reason is that these payloads only use an unique API call (*WinExec*), and before invoking it, one of the following APIs is kept as a return address in the stack:

- `ExitThread`
- `ExitProcess`
- `SetUnhandledExceptionFilter`

The idea is to terminate cleanly the compromised process once the selected command is executed.

When *Shelia* verified which address *WinExec* was invoked from, it determined that it was from a correct address due to `kernell32.dll`, the library where the previously referred functions reside, is loaded in a executable region. So *Shelia* missed this attack.

The attack detection engine was consequently improved to deal with the described situation. The patch was trivial and after this enhancement, *Shelia* was able to detect all of these payloads for each vulnerability tested.

win32_exec. Shelia was able to log which command was to be executed. Moreover, if containment policies were active, the command execution was rejected.

win32_adduser. This payload can be considered equivalent to the previous one, with the particularity of executing the shell command: net adduser.

win32_bind. This payload spawns a command shell listening for commands from the attacker. If no containment measures were active, Sheila was able to log the conversation maintained between the attacker and the command shell. No matter if SSL encryption was activated during network communication since data is caught unencrypted at application level.

win32_downloadexec. Shelia was able to identify the payload downloaded and keep it in a special folder. If containment measures were active, the payload execution was rejected.

2.8 Performance

We ran a set of micro-benchmarks to determine the (worst-case) performance hit on system-call invocations, and macro-benchmarks to determine the performance impact on real-world software. We used the ApacheBench benchmarking tool from the Apache HTTPServer project for the macro-benchmarks. Our results show that Shelia imposes a small performance overhead on average for Apache 2.0.55 (3% reduction in request processing capability). The benchmark was run on a 2.8GHz Pentium IV laptop with 512MB of RAM. Our experiments should be taken as indicative, rather than exhaustive, because for space reasons we only include the most pertinent experiments in this deliverable.

2.8.1 Micro-benchmarks

We wrote a C program that uses the *QueryPerformanceCounter* function to measure the system time to invoke different Win32 APIs LIMIT times in a tight loop (LIMIT=1000).

We then divided the total time by LIMIT to get the mean execution time per invocation. We collected such values from 10 runs, and averaged the median 8 values for each system call. Figure 2.3 shows the results for individual microbenchmarks to assess upper-bound performance overhead for some potentially dangerous system calls. The overhead settles in the range 1-2 μ -seconds per 4-6 μ -seconds systemcall execution time.

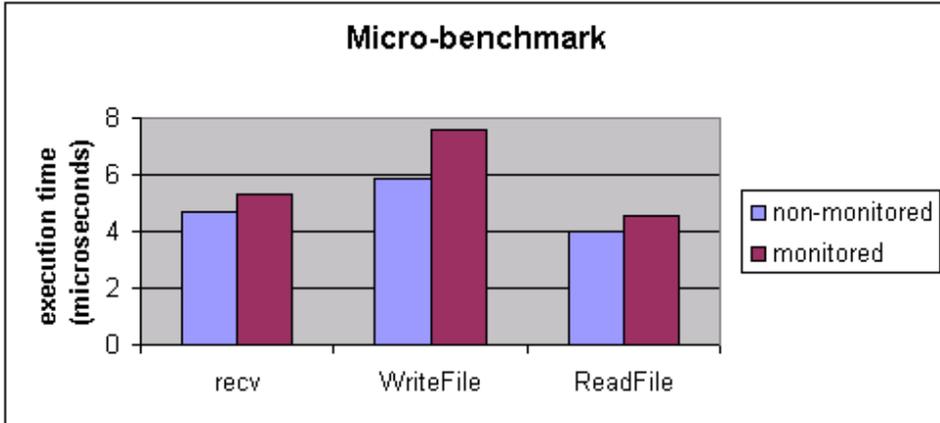


Figure 2.3: Micro-benchmark test results

2.8.2 Macro-benchmarks

We chose the ApacheBench benchmarking suite for the Apache HTTP server project as a realistic benchmark for evaluating the performance impact of Shelia. Apache is ideal for this purpose owing to its wide-spread adoption, and to the fact that it exercises many of the hooked functions such as *connect*, *send*, *CreateProcess*, *ReadFile*, and *LoadLibrary*. We ran the ab tool from ApacheBench with the parameters: `-n 10000 -C 1000 http://localhost/70KB.bin` to simulate 1000 concurrent clients making a total of 10,000 requests for a 70KB file through the loopback network device (i.e., on the same host as the web server, to avoid network latency-induced perturbations to our results). We collected and averaged the results for 5 runs of ab for each server configuration.

	Apache	Apache/Shelia	Slowdown
req/s	350,55 ± 2,05	340,32 ± 1.19	2.92% ± 0,90
thruput (kB/s)	24104,53 ± 141,14	23401,23 ± 81,57	2.92% ± 0,90

Table 2.1: Shelia macro-benchmark: ApacheBench.

Table 2.1 and figure 2.4 show the results of our macro-benchmark tests. The Apache server suffers a modest 3% decrease in request-handling capacity while running under full monitored environment, as compared to running on a stock system. Moreover, the standard deviation indicates that this decrease is not even very significant.

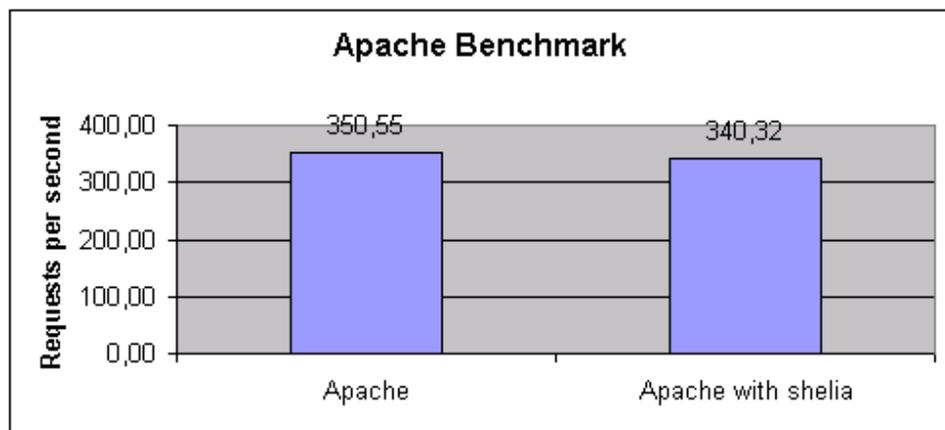


Figure 2.4: Apache Benchmark. Request-handling capacity decrease measured.

Chapter 3

Eudaemon: a good spirit

‘Greeks divided daemons into good and evil categories: eudaemons (also called kalodaemons) and kakodaemons, respectively. Eudaemons resembled the Abrahamic idea of the guardian angel; they watched over mortals to help keep them out of trouble. (Thus eudaemonia, originally the state of having a eudaemon, came to mean “well-being” or “happiness”.)’ [Wikipedia]

Sophisticated high-interaction honeypots like Argos [29], Minos [10], and Vigilante [23] all use dynamic taint analysis, as pioneered by Denning [14] and evolved by Newsome and Song to capture zero-day attacks in TaintCheck [26]. In essence, data originating from suspicious origins (e.g., network data) are tagged “tainted”, and an alert is raised when they are used to divert control flow, or when they are executed. In practice, taint analysis is performed using emulators or binary re-writing tools and incurs an overhead ranging from one to several orders of magnitude. As such, honeypots are ill-suited for full-time deployment on end-user production systems. Additionally, current honeypots have several fundamental disadvantages that severely limit their usefulness [18, 41]:

1. Honeypot avoidance: an attacker may create a hitlist containing all hosts that are not honeypots and attack only those machines.
2. Configuration divergence: the configuration of honeypots often does not match exactly the configuration of production machines. For instance, users may have installed different versions of the software, plugins, or additional programs. Honeypots only reflect a limited set of configurations. Indeed, high interaction honeypots typically have a single configuration.
3. Management overhead: honeypots require administrators to manage at least two installations: one for the real systems, and one for the honeypot.

4. Limited coverage: even if a honeypot covers a sizable number of IP addresses, it may take a while before it gets infected. This is especially true if the honeypot only covers dark IP space. Moreover, the address space that is covered is limited by the amount of traffic that can be handled by a single honeypot.
5. Server-side protection: most honeypots mimic servers, by sitting in dark IP space and waiting for attackers to contact them. Unfortunately, the trend is for attackers to move away from the servers in favour of client-side attacks [15, 31].

Other intrusion detection methods that do not rely on taint analysis and perform better than it do exist, but suffer from other problems. We saw earlier that *Shelia* uses a method that checks whether or not a call came from the data segment. That works quite well, except when multiple stepping stones (library calls) were used to get to the sensitive operation. In that case, the call came from an executable region. We saw this problem creeping up in the evaluation section. Moreover, *Shelia* suffers from most of the other problems listed above. Plus it detects attacks when they have already started executing instruction, which is potentially harmful in its own right.

Measures like StackGuard [11], and address/instruction set randomisation [7, 17] are cheap, but can be overcome [32] and do not enable the generation of any type of “vaccine” for the exploited fault. Replaying identified attacks against high-interaction honeypots has been suggested [23, 36] to address the latter issue, but successful replaying remains a subject of research in the presence of challenge/response authentication [12, 20]. Moreover, heavily instrumented applications or machines that serve as replay targets for many alerts do not scale easily.

To solve the honeypot problems mentioned above, without sacrificing the generation of valuable data about the attack, we propose to turn end-user hosts into heavily instrumented honeypots. This can be achieved by transparently switching any application between native and intensely instrumented execution, whenever desired and in a timely manner. We call this architecture *Eudaemon*¹. Previous work in this area proposed selective protection of a particular *segment* of an application [33, 22]. Running the segment in instrumented mode provides a means to generate patches that ‘fix’ the faults. However, these solutions are dependent on a detector that will initially identify the attacks. We therefore claim that they are complementary to *Eudaemon*.

As an alternative, Ho et al. investigated ways to speed up taint analysis so as to make it deployable as a full-time solution on production machines [4].

¹The Eudaemon work was published at the ACM SIGOPS EUROSYS in April 2008 [28].

Their solution is based on a virtual machine (VM), that transfers execution to a modified Qemu [6] emulator, whenever tainted data are read into the system or processed. They achieve much better performance than other systems providing system-wide protection, but the slow-down is still significant (a factor 2 on average). In addition, they require the installation of a modified Xen hypervisor on the machine which in practice hinders its deployment on the majority of home users' PCs. Finally, while full-system protection is attractive as it also catches attacks on the kernel, the downside is that it becomes harder to provide fine-grained analysis of the actual program under attack.

Ideally, one would make every host operate under heavy-weight instrumentation constantly so as to provide full-time safety. Unfortunately, as we have seen, doing so is impractical (at least in the foreseeable future) due to the associated overhead which would likely result in a reduction in user productivity. On the other hand, we propose that it may be possible to explicitly switch to heavily instrumented 'honeypot mode' under certain conditions, provided the conditions are such that they strike a balance between increased protection and performance. In the remainder of this section, we sketch two such scenarios: idle-time honeypots, and honey-on-demand.

Idle-time honeypots

Studies suggest that PCs tend to be idle more than 85% of the time [16]. This refers to both idleness due to lack of user interaction (idle desktop), and idleness in terms of processing (idle CPU). Client machines display both types, but the former presents an interesting opportunity, and can serve as the condition that triggers the switch to honeypot mode. Much like a screen-saver protecting the screen from damage while the user is away, turning a machine to a honeypot protects running processes (e.g., instant messengers, p2p programs, system services etc) from attacks such as buffer overflows, code injection etc. While acting as a honeypot the machine behaves exactly as it did before, with the sole difference being a reduction in processing speed.

If at any time, any host can be a honeypot, the rules of the game for the attackers change significantly. For instance, they can no longer harvest a set of IP addresses in advance, because what appears to be a suitable target now may be a heavily instrumented honeypot by the time you attack it. As long as some machines in the set run as honeypot, the attacker risks being exposed. As a result, important classes of attack are rendered obsolete, and honeypot problems 1-4 from above are resolved.

Honey-on-demand

An alternative application of *Eudaemon* is sometimes popularly referred to as ‘the big red button’, i.e., a button that users may press when they are about to access an unknown and possibly suspicious website, or when they open attachments, view images or movies, etc. Pressing the button will make the application run in honeypot mode, heavily instrumented and safe against client-side exploits.

As it may often seem ill-advised to depend on the user for making such decisions², we stress that the ‘big red button’ is a metaphor. It represents a generic interface for determining which application should be protected when. Besides end users, the interface could be used by applications. For instance, mail readers could demand to be run in emulation mode when opening an email classified as spam, or from an unknown sender. Also, faster but less accurate intrusion detection systems or access control systems could trigger a switch to honeypot mode in the event of an anomaly. Alternatively, other mechanisms, such as whitelists of network addresses could be used to decide whether a web browser should switch to emulated execution.

Using *Eudaemon* in the above manner helps us tackle the last and potentially most important issue with current honeypots. Since 2003, client-side attacks are increasingly common. Hackers take over client machines and group them into botnets that are subsequently used for unwanted and illegal activities, such as spamming, on-line fraud, distributed denial of service (DDoS) attacks, and harvesting of passwords and credit cards. Such attacks are not caught by most current honeypots and client honeypots are much less common, and for the few that do exist (e.g., [38, 24]), the other problems remain.

Finally, *Eudaemon* may be used for servers. Often, when a vulnerability is announced, there is not yet a patch available. Even if there is a patch available, administrators may be reluctant to apply it right away. If the server is not too heavily loaded, *Eudaemon* may be used to run the server in safe mode, thus buying precious time until the patch can be applied. Such usage escapes the honeypot and client-side exploits domain, and enters the area of intrusion prevention.

Contribution: Eudaemon

In this chapter, we present the design, implementation and evaluation of *Eudaemon*, a ‘good spirit’ capable of temporarily possessing unmodified applications at runtime to protect them from evil. The contribution of this chapter is a novel idea for applying honeypots, with a wide range of possible applications. Our focus is primarily on the *techniques* for possession,

²On the other hand, similar principles are used extensively in a modern OS like Windows Vista

protection, and release, rather than on the applications that may make use of them. In addition we explain in detail how such a switch to and from honeypot mode works in a modern operating system.

In a nutshell, when *Eudaemon* receives the order to possess a process, it attaches to the process in such a way that we can observe and control the execution of the target process, and examine and change its core image and registers. Most modern OSes have functionality for doing so (for instance, UNIX provides the `ptrace` system call for this purpose, while Windows XP allows for the creation of threads in target processes). We temporarily pause the execution of the victim process, save its processor state, and inject a small amount of shellcode in its address space. The shellcode calls a modified version of an open source emulator which is linked in as a library. The emulator is started with the processor state that was previously saved. From that point onwards, execution of the process code continues, except that the emulator provides full-blown taint analysis, and raises an alert whenever data with suspicious origins (e.g., the network) is used in a way that violates the security policy. When *Eudaemon* receives the order to release the process, it halts the process temporarily, removes itself from the process and resumes the process in native mode. In other words, network applications (e.g., peer-to-peer or FTP download systems), besides being inactive for a few milliseconds, are not interrupted for the possession period.

The remainder of this chapter is organised as follows. In Section 3.1 we place our work in the context of related work. Section 3.2 presents an overview of the system's design. Implementation details are given in section 3.3. Section 3.4 evaluates performance, and conclusions are drawn in Section 4.

3.1 Related Work

Taint analysis is used by many projects such as TaintCheck [26], Minos [10], Vigilante [23], and Argos [29]. Most of the existing work assumes deployment on dedicated honeypots. This is mainly due to performance reasons. Likewise, client-side honeypots tend to be dedicated machines also [38, 24]. As a result, these techniques suffer from most of the problems identified in Section 3.

A interesting exception includes the work on speeding up taint analysis by switching between a fast VM and a heavily instrumented emulator by Ho et al. discussed earlier [4]. One drawback of the method (besides an overhead that is still considerable compared to native execution) is that it can only be installed by, say, home users willing to completely reconfigure their systems to run on a hypervisor.

In contrast, we deal with performance penalties by running in slow mode on demand. In essence, we slice up program execution in the temporal do-

main. A different way of slicing (proposed by a group at Columbia University) is known as application communities [21]: assuming a software monoculture, each node running a particular application executes part of it in emulated mode. In other words, applications are sliced up in the spatial domain and a distributed detector is needed to cover the full application. *Eudaemon* directly benefits individual installations without relying on a monoculture. In practice, the OS used by communities tends to be uniform, but variation exists in applications, due to plug-ins, customisations and other extensions.

In a later paper, the same group at Columbia employs selective emulation to provide self-healing software by means of error virtualisation [22]. Again slicing is mostly in the spatial domain. As far as we are aware, neither of these projects supports taint analysis. Indeed, it seems that for meaningful taint analysis, the tainted data must be tracked through all functions and, thus, *selective* emulation may be more problematic. At any rate, as we mentioned earlier, the *Eudaemon* technique is complementary to [22] and could be used as an attack detector.

Another interesting way of coping with the slowdown (and indeed, a way of slicing in the temporal domain for servers) is known as shadow honeypots [5]. A fast method is used to crudely classify certain traffic as suspect with few false negative but some false positives. Such traffic is then handled by a (slower) high-interaction honeypot (e.g., Argos). Tuning the classifier is delicate as false positives may overload the server. In addition, shadow honeypots suffer from the problems of configuration and management overhead identified in Section 3.

Rather than incurring a slow-down at the end users' machine, many projects have investigated means of protection for *services* running on user machines by way of signatures and filters. Such projects include advanced signature generators (e.g., Vigilante, VSEF, and Prospector [23, 8, 34]), firewalls [40], and intrusion prevention systems on the network card [13].

For instance, Brumley et al. propose vulnerability-based signatures [8] that match a set of inputs (strings) satisfying a vulnerability condition (a specification of a particular type of program bug) in the program. When furnished with the program, the exploit string, a vulnerability condition, and the execution trace, the analysis creates the vulnerability signature for different representations, Turing machine, symbolic constraint, and regular expression signatures.

Packet Vaccine [37] detects anomalous payloads, e.g., a byte sequence resembling a jump address, and randomises it. Thus, exploits trigger an exception in a vulnerable program. Next, it finds out information about the attack (e.g., the corrupted pointer and its location in the packet), and generates a signature, which for instance can be based on determination of the roles played by individual bytes. To determine this, Packet Vaccine scrambles individual bytes in an effort to identify the essential inputs.

Vigilante relies on the possibility to replay attacks in order to generate Self-Certifying Alerts (SCAs), an extremely powerful concept that allows the recipient of an alert to check whether the service is indeed at risk [23]. If so, it may automatically generate a signature. The false positives rate seems to be zero.

As mentioned earlier, the signature generators for the above projects may be capable of handling zero-day attacks, but they produce them by means of dedicated server-based honeypots. Hence, they do not cater to zero-day attacks on client applications. To some extent the problem of zero-day attacks also holds for virus scanners [35], except that modern virus scanners do emulate some of the data (e.g., some email attachments) received from the network. However, remote exploits are typically beyond their capabilities.

We should mention that we currently have a few signature generators for our emulator (known as Sweetbait [29] and Prospector [34], respectively). As part of our future work, we intend to use different types of signature generator (some with light-weight instrumentation, and other with heavy-weight analysis), so that different users may apply different forms of analysis to the same attack, hopefully yielding a more complete picture of the attack.

Protection mechanisms such as StackGuard [11], PointGuard [9], and address space and instruction set randomisation [7, 17] protect against certain classes of attack, but are unable to generate much analysis information about the attack, let alone generate signatures.

Many groups have tried to use such fast detection methods like address space randomisation and perform more detailed instrumentation on a different host by replaying the attack [36]. In our opinion, replaying is still difficult due to challenge/response authentication, although promising results have been attained [12, 25]. More importantly, the heavily instrumented machines that perform the analysis may become a bottleneck if many attacks are detected. *Eudaemon* inherently scales because it employs the users' machines.

Process hijacking is a common technique in the black-hat community [2, 30]. By injecting code into live processes, such attacks are hard to detect, as no separate process is created and no attack can be found at the file-system level.

To conclude this section, in 2005 Butler Lampson proposed to partition the world into two zones: green (safe) and red (unaccountable) [19] and use a VM to isolate the two parts. While more work is clearly needed in this area, we believe *Eudaemon* might be a step toward having the two zones while maintaining an integrated view on the applications.

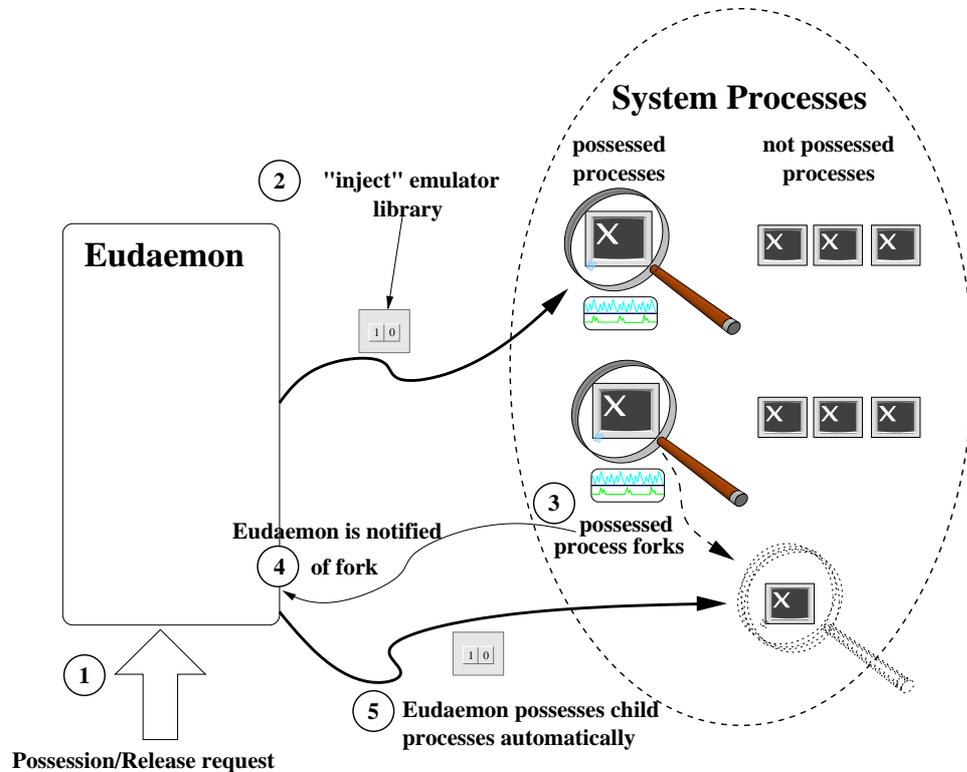


Figure 3.1: Eudaemon Overview

3.2 Design

Eudaemon has been partially inspired by techniques used by hackers and debuggers alike to attach to running applications, instead of requiring them to be loaded from within the controlled debugger context. We use similar techniques to hijack or *possess* a process transparently with the goal of heavily instrumenting unmodified binaries to protect them against remote exploit attempts.

Eudaemon has been designed to run as a system service, where requests to possess or release applications can be made. The terms *possession* and *release* will be used to describe the act of switching a process to emulated and native execution respectively. A high level overview of the system is shown in Figure 3.1. Requests to possess or release an application can be issued based on any criterion, such as an explicit user request, or as a result of persisting inactivity at the host, as mentioned in the introduction.

After receiving a request (①), *Eudaemon* immediately attempts to attach to the target process to force it to run in an emulator (②). The complexity of the procedure varies depending on the platform implementation, but most modern operating systems do support (system) calls that

implement the desired functionality (e.g. Linux, BSD, and Windows XP). Attaching to a process can be performed using its process identifier (PID) alone. When threading is used, the thread identifier (TID) can be used instead.

For safety, the operating system ensures that only a program running under the same or super-user credentials is able to attach to a given process. This scheme guarantees that users cannot possess or release processes they do not own.

When an emulated process spawns new processes (③), we can also request the automatic possession of its children to enable *Eudaemon* to protect an application consisting of multiple processes (⑤). We emphasise that even in this case a program need not be *Eudaemon*-enabled in any way. The emulator library which marshals execution makes sure to notify the system on the creation of new processes (④). Threads can be handled internally, since all of them share the same “possessed” address space, and the emulator can proxy new thread requests.

3.2.1 Process Possession

Switching to emulated execution (possession), is accomplished by *injecting* code to perform this task within the target process space. For threaded applications it is sufficient to inject the code once, since all threads lie in the same address space. Nevertheless, the amount of code required to perform such a complex job can be significant. What this implies is that the costs of copying or injecting the emulator code within a process, could compromise the transparency of the system.

We overcome such an eventuality by making the emulator a dynamic shared library (known as DSO or DLL, depending on the platform). Libraries impose some restrictions on the included code, but on the other hand make code reusing simple and efficient. When a DLL is loaded by multiple processes, it is actually loaded once in system memory, and only mapped in each process’s address space. As operating systems allow libraries to be loaded either at runtime, or *a priori* for every process, we have some freedom in how we inject the emulator code efficiently and in a way that will scale even when multiple targets are possessed. For more details on loading the emulator in a process, we refer the reader to Section 3.3.

After loading the necessary code in the target process, we still need to activate it, and supply the required state so that execution can resume virtually undisturbed. Acquiring a target’s execution state is commonly performed by debuggers and we use the same technique. As we will explain in detail shortly, supplying such state to the newly injected code in the target, is more involved and requires that we protect the integrity of the target. Phrased differently, our code shares the target process’ address space, and we need to isolate its memory from the process’s memory.

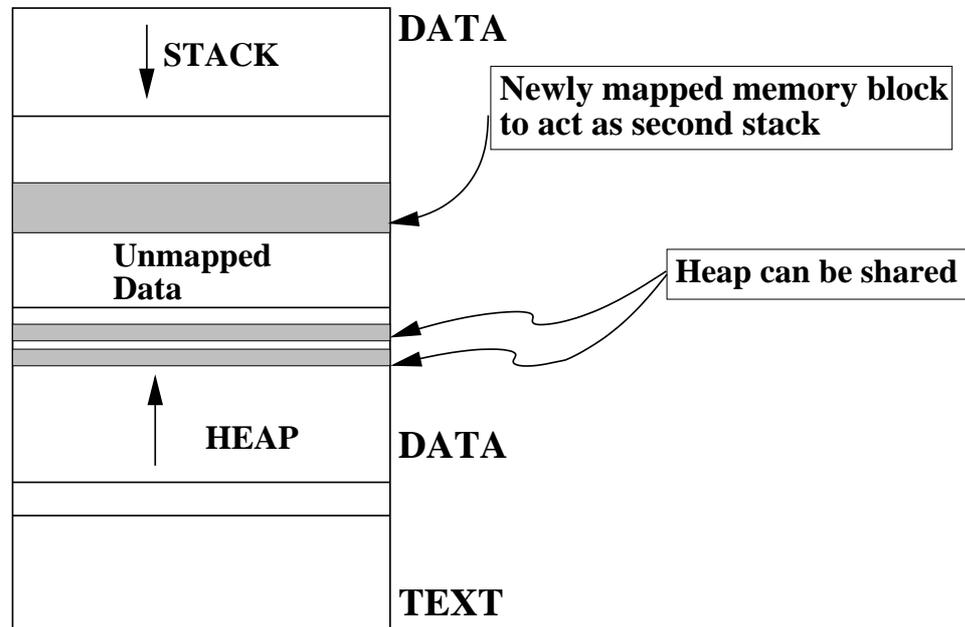


Figure 3.2: Process Memory Layout

Figure 3.2 shows the typical layout of a process' memory. Application code is loaded in what is called the *text* segment, and it is protected by being marked read-only. Loaded libraries, even though not shown in the figure, also have their corresponding code segments protected by being read-only, and as such our code is implicitly protected. Process data are stored mainly in two areas: the heap and the stack. Heap size is dynamic, and grows towards larger addresses, while stack is usually fixed, and is used as a LIFO queue. The stack grows towards lower addresses.

Every thread of execution has its own stack, which in architectures like the x86 is addressed using special CPU registers and implicitly updated by special instructions. As a result, executing our code using the same stack as the process code would lead to severe inconsistencies in the stack. In contrast, heap memory is larger and allocated objects are referenced explicitly by holding memory pointers. This permits us to share the heap for any objects we need to allocate. In both cases, however, we cannot rule out the possibility that the program will access data owned by the library either as a result of an error, or as part of a malicious attempt to thwart its proper operation. We handle data protection through the emulator, which we describe in more detail in Section 3.2.3.

To call safely the code we have already injected, we first map a memory segment in the target process that will serve as a stack for the emulator. This way we ensure that both our code and the process' code can be run in parallel without interfering with each other. The way this is accomplished

depends on the underlying system. For example, some systems allow a process to reserve a memory segment in another process, while on others we are forced to inject a piece of short lived code to perform the allocation.

In the latter case, we need to choose carefully the location where to place the short lived *activation* code, so as to not compromise the target's integrity. One possible solution is to choose an area in the target's text memory space, save it, and then overwrite it with the activation code. When the activation code completes, the original code can be restored. However, this process requires pausing all threads in order to guarantee that the location will not be used while the activation code resides there.

Ideally, we would like to avoid such overhead and prefer to inject the code in a location that we know is no longer in use. In practice, the binary's header that resides in the text segment is often a feasible location. Certain bytes in the header are used only when the executable is loaded by the system in memory. Usage of executable header memory to run code has been demonstrated before by virus writers to inject parasitic code in running programs [3]. A more extreme solution is to use the space left by compilers between the functions of an application for performance reasons [1], but we have not explored such a course in our implementation.

To activate the library, we use a part of the newly allocated stack to store the state we extracted when we attached to the target, and detach from the process. Finally, the activation code calls a function in our library which takes over the execution of the process.

3.2.2 Process Release

To return a process back to native execution the emulator needs to be notified to clean up and export the state of the process, as it would have been if the process has been running natively all the time. Delivering such a notification could be performed by various means, but to preserve semantics similar to those of possession we chose once more to inject *deactivation* code into the process. The code simply performs a call within the library to deliver the notification.

If the call succeeds then *Eudaemon* needs to wait until the emulator exits, and control is returned to the activation code that was injected during possession. The remainder of that code will notify *Eudaemon* that the process can be switched back to native execution, and if necessary also release the allocated second stack. To complete the switch, *Eudaemon* reads the state of the process as exported by the library, and reinstates it as the process' native state.

3.2.3 Emulator Library

The emulator library is decoupled from *Eudaemon*, so that it can be transparently replaced without affecting the system's operation. As long as the library adheres to *Eudaemon*'s predefined emulator interface, and the library itself does not compromise the process, any implementation that shields the process against attacks can be used. We now describe at a high level the required interface for a library to be used in *Eudaemon*, and also present the criteria that need to be obeyed in the remainder of this section. The exact library calls will be listed in Section 3.3. From a high-level perspective, the desired interface consists of three functions:

- We need a function to check that the library is not already in control of the target process in order to handle requests to possess a process that is already possessed. To avoid possible conflicts the state of the library (active/inactive) is exported via such a call.
- A function is needed to initialise and give control over the process to the library. The function represents the entry point of the library, where control is redirected after setting up memory and process state. It should neither fail nor return, unless there is an error in the program itself or the library was notified to exit.
- The final function we require is one that signals the emulator library to relinquish control of the process, and return to the caller. This call need not be synchronous, in the sense that the library does not need to terminate immediately. *Eudaemon* will wait for the process to complete the switch to native execution.

A more important aspect of the library is that it should protect itself from unintentional or malevolent access of critical data. As we briefly mentioned earlier, a program could access library data in stack or heap, and in that way compromise the mechanism that is supposed to be protecting the application. To guard against such a possibility we adopt a memory protection method very similar to the one used in the x86 CPU architecture (see Section 3.3.1.4 for details).

3.3 Implementation

We completed an implementation of *Eudaemon* on Linux. We also completed most of the possession and release functionality for Windows, while work on the required library-based emulator is in early stages. Due to size restrictions, in the remainder of this section, we only discuss in detail the Linux implementation of the main components of our design: (i) a process emulator that implements taint analysis, and (ii) *Eudaemon* possession and release.

3.3.1 SEAL: A Secure Emulator Library

SEAL is a secure, x86-based user-space process emulator implemented as a library. It is based on Argos and employs the same dynamic taint analysis [29]. In a nutshell, Argos is a whole-system emulator and consists of a virtual CPU, NICs, video card, etc. It marks all data arriving on its virtual NICs as tainted and tracks them throughout their lifetime in the system. Argos raises an alert when tainted data is used in illegal ways (e.g., when it is executed).

We modified Argos in the following ways. First, we do not desire whole-system emulation, so we ported the dynamic taint analysis functionality to a user-space emulator. As Argos shares its code base with Qemu [6], which includes a user-space emulator, doing so was straightforward. Second, a user-space emulator has no notion of virtual NICs, so we had to modify the tagging mechanism. For instance, SEAL tags bytes when they are read from sockets (and certain other descriptors). Third, as the original process and SEAL share the same address space, we had to protect data used by SEAL from being clobbered by the process. Fourth, we packaged SEAL as a library with a succinct interface. We now discuss these issues and the general operation of SEAL in detail.

3.3.1.1 Tagging Data

Processes read data by means of the `read` system call which is used for sockets, files and pipes. To distinguish suspect data from harmless input, we introduce a 64KB bitmap (a bit for each one of the possible 2^{16} descriptors in Linux) that marks certain descriptors as tainted. Calls to `read` result in data tagging only if a tainted descriptor was used. We now describe how we monitor system calls to taint descriptors. First, `socket()` and `accept()` both create descriptors for network communication. As network data is suspect, the descriptors are marked tainted. Second, `open()` returns a descriptor for file access. Normally, we ignore this call, but users are allowed to mark certain directories as *unsafe* to capture exploits in files. Consider a malicious image in an attachment that triggers a vulnerability in the viewer. SEAL scans the path name provided to `open`, and taints the descriptor if the file is in a directory marked *unsafe* (e.g., `/tmp`, or `/home/...`). Third, the `pipe()` call creates a pair of descriptors for inter-process communication. SEAL considers input from another process as unsafe, since it is of unknown origin, and taints both descriptors. Finally, `dup()` and `dup2()` create a copy of a descriptor. If the original descriptor is tainted, we also taint the copy.

Besides the `read()` system call, programs can access input by means of message passing and memory sharing. Messages can be exchanged either over a network socket, or a message queue. In both cases, we can trivially monitor the message receiving system calls to taint incoming data. Han-

dling shared memory is more difficult. Programs may either map files into their address spaces, or share memory pages with other programs. Simply tainting the memory is not sufficient, because it would miss updates made by other processes. We therefore included a sticky flag for every tainted page. Asserting this flag ensures that the page will be always considered tainted ignoring all writes performed by the process, until it is unmapped or not shared anymore.

3.3.1.2 Tracking Tainted Data

Data items tagged as *tainted* are tracked during execution. Tracking is achieved by instrumenting the guest's instructions to propagate tags. For example, arithmetic operations like ADD and SUB, are instrumented to taint their result, whenever they are used with a tainted operand. In a similar way, MMU operations such as load and store, copy tag values between registers and memory. The tagging granularity is variable. Memory data are tagged per byte, while a single tag is assigned to each of the 8 CPU registers. MMX registers are treated as memory, and have byte granularity as well. As a result, every byte that depends on network data can be monitored.

Tags in SEAL are accessed through a one-level page table. We partition memory space in pages, and only when data belonging to a memory page are tainted, tag space for that page is allocated and the corresponding tags asserted. The page table contains pointers to structures actually containing the tags for each page, where a tag can either be a single bit, or a byte. While it would have been faster to use a one-dimensional tag array, we wanted to keep the memory footprint of the emulator as small as possible, especially since SEAL and user application share the same address space. In addition, by aligning the dynamically allocated blocks of tags on addresses that are multiples of four, the least significant bits of page tables entries are unused, and can be used to track inexpensively the sticky page flag mentioned above.

When a typical Linux process is running SEAL using single bit tags, the amount of memory X (in MB) that can be used by the process can be expressed as: $X + (X/8) + 4 < 3072$ (the maximum addressable virtual memory being 3 GB or 3072 MB, the page table taking up to 3 MB, and 1 MB taken by library code and statically allocated data). So, a process under SEAL can use up to 2727 MB of the virtual address space, reducing its maximum available memory by about 9.64%. At runtime, the actual memory footprint of the library depends on application behaviour, and the amount of tainted data. We can calculate a 12.5% upper boundary for the memory overhead imposed by the library, if we assume all process data are tainted and a single bit is used for each byte.

3.3.1.3 Attack Detection

Most exploits attempt to redirect control to shellcode provided by an attacker, or to code that is already available (libc). They do this either by loading an attacker supplied value on the instruction pointer (EIP), or by injecting instructions within a program's control flow. On x86 CPUs EIP is manipulated using one of the *call*, *jmp*, and *ret* instructions. SEAL monitors these instructions, and checks that none of them is used with tainted arguments, or results in EIP pointing to tainted data. Even in the case where EIP is not directly pointed to a tainted location, "walking in" an area with tainted code will eventually cause an alert since attackers are bound to use a checked instruction (such as *jmp*, *call*, or *ret*). In other words, SEAL is able to detect most overwriting and code injecting exploits.

After an attack is detected, SEAL generates an alert and logs the state of the emulator to persistent storage. It scans the victim process's memory and logs all locations that have been marked tainted, as well as the virtual CPU's registers, and the type of the offending instruction. It also collects information (like pid, name, and DLLs) about the victim application. The logs are subsequently used by signature generators to create anti-measures. Signature generation in Argos/SEAL is beyond the scope of this paper. Interested readers are referred to Sweetbait [29] and Prospector [34].

3.3.1.4 Protecting SEAL Data

As application and SEAL reside in the same address space, we need to protect emulator data against malicious or accidental accesses by the application. As mentioned earlier, our solution resembles memory protection in x86 architectures. The x86 CPU contains a hardware memory management unit (MMU) that partitions the linear physical memory space into pages of virtual memory space. The MMU is responsible both for translating a virtual address to a physical one, as well as enforcing a page protection mechanism. This way every process is assigned each own virtual address space isolating it from other processes, and protecting kernel space from the processes.

We adopt the same principle by using a virtual MMU that enforces page level protection. SEAL instruments all memory accesses in the application's code to go through the virtual MMU, where they are validated to make sure that library owned memory is not accessed. Every page allocated by the library is marked with a flag that allows the virtual MMU to perform the validation. The structure that these flags are stored in is of small importance; a reasonable choice in our case was to use one of the extra bits in the page table described in Section 3.3.1.2.

Keeping track of protected memory pages is straightforward. It only requires that on allocation and release of heap memory the library updates the

corresponding bits. The virtual MMU can also be used to protect the stack, global library data, and library read-only data to defend against information leakage that could be exploited by attackers. Obviously, it protects its own bitmap and data, while the code is protected in the same way as all other code.

3.3.1.5 Checking System Calls

Monitoring the use of tainted data in critical operations is the same as in the whole-system emulator. However, being in user-space offers us the chance to expand operations that are monitored to include certain system calls. In theory, we could apply policies concerning tainted arguments to all system calls, but in practice it makes sense primarily for the *exec()* system call which executes a file by replacing the image of the current process with the one in the file. It has been frequently exploited by overwriting the arguments to load arbitrary executables. By checking the arguments for tags, SEAL shields programs against such attacks.

3.3.1.6 Signal Handling

SEAL handles signals transparently to the application. Upon receiving control of a process, original signal handlers are replaced with the emulator's handler. This single signal reception point queues arriving signals that will be processed at a later time. System calls used to update signal handlers and masks, are also intercepted to keep track of the process's signal related behaviour.

Such an approach is necessary to ensure that native code is never called directly, but to allow also switching to emulation mode while executing a signal handler at the target. To clarify this point, we will briefly describe how the Linux kernel handles signals. Upon signal delivery, a new temporary execution context is created by the kernel for the handler to execute, and the previous context is saved in user-space. Before relinquishing control to the signal handler, which runs in user-space, a call to *sigreturn()* is injected in the temporary handler context. This system call serves the sole purpose of returning control to the kernel, so that the original execution context can be restored. When SEAL is in place, it imitates the kernel. As a result, if the emulator receives control while a signal handler is executing, it is still able to switch to the process's original execution context in emulation mode by intercepting *sigreturn()*.

3.3.1.7 Eudaemon Support

The SEAL user-space emulator as described so far, can be used to run applications securely, but cannot be used with *Eudaemon* yet. To enable the transition of a process from native to emulated execution we need further

extensions. Primarily, SEAL needs to be in a form which can be dynamically included in any process. Dynamic shared libraries or DLLs provide exactly that. Compiling SEAL as a dynamic shared library is trivial, but it requires a simple interface to interconnect with *Eudaemon*. We use the following as a basic interface for interconnection with *Eudaemon*:

- *bool seal_isrunning()*; this function receives no arguments. It returns a boolean value that specifies whether the emulator is active at the moment the function was called.
- *void seal_initandrun(struct cpu_state *st)*. This is the library's main entry point. It initialises the emulator with the snatched process state such as register values, MMX, and floating-point state, and commences emulation.
- *bool seal_stop()*; this function requests that the emulator stops, and consequently that *seal_initandrun()* returns. It returns *true* on success and *false* if SEAL is not actually running. Calling this function does not cause the emulator to exit immediately. Instead it waits until the virtual CPU reaches a state that is safe to return.

Finally, the *exec()* system call also requires modification. Compiling SEAL as a library means that if the current process image is replaced with a different executable by *exec()*, we have to re-attach and switch it to emulation mode, or let the newly called binary execute natively. By default we use the latter option. To support the former, we permit *exec()* to signal *Eudaemon* of the event, so that the new process can be forced into emulation mode once again.

3.3.2 Possession And Release

Process possession and release are two distinct operations that are independent in the sense that no state needs to be preserved between the two. In other words, a possessed process holds all the information needed for its release. The only prerequisite for these operations is that the emulator library is present in the target process's address space.

The finer details of injecting the library in the target process, as well as activating and deactivating it are in some cases very dependent on the underlying OS platform. In the remainder of this section, we elaborate on the implementation details of *Eudaemon* on Linux.

We use the shared library pre-loading mechanism in Linux to transparently load the emulator library in the address space of every process. In detail, Linux and other Unix based systems support the pre-loading of dynamic shared libraries in applications using dynamic linking. This is accomplished by either defining the environment variable *LD_PRELOAD* to

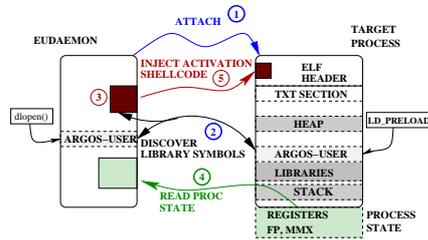


Figure 3.3: Process possession: phase 1

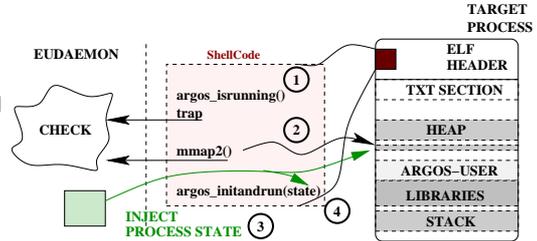


Figure 3.4: Process possession: phase 2

include the desired library, or by including it in a configuration file (like */etc/ld.so.preload*).

Eudaemon employs the Unix system call *ptrace()*, which was originally intended mainly for debugging purposes. Much like a debugger, we use *ptrace()* to achieve possession and release without process and OS cooperation. In summary, *ptrace()* allows one process to attach itself to another, assuming the role of its parent. The target is stopped and the attaching process receives control over it. The attaching process is then able to read the target's state, such as register values, floating point (FP) and MMX state, as well as memory data. It is also able to resume execution of the target process, while receiving notification of events such as system call execution and signal reception. This allows *Eudaemon* to access process state, and to inject the instructions needed to perform the switch from native execution to emulation and vice versa.

3.3.2.1 Process Possession

The possession operation can be logically split in two phases. The first phase is shown in Figure 3.3 and consists of the following steps: (1) attach to target process; (2) discover necessary emulator library symbols in the target; (3) modify activation shellcode using the symbol addresses acquired during step 2. Each of these steps will be explained in more detail below.

To possess a process we first attach to it, and wait until the target is effectively stopped by the OS. Subsequently, we look up the target's memory mappings to find out the location of the emulator library in its address space. We accomplish this by looking up */proc/[pid]/maps*, where *[pid]* is the target PID. This is a file under the special *proc* filesystem, and contains a description of the memory mappings used by each process. Figure 3.5 shows the contents of such a file. Every line of this file corresponds to a memory mapping and provides information on its address range, protection bits, size and source filename, if applicable. We are thus able to locate the address where the emulator library was loaded in the target, as well as in

Eudaemon itself. Observe that `libseal` is listed twice in the file. The reason for this is that `bss` is also listed.

With this information, we can at runtime look up any emulator symbol in the target. We accomplish this by also loading the emulator dynamic shared library in *Eudaemon*, using dynamic loading and linking, and calculating the offset of the symbol from the beginning of the dynamic shared library. The offset of the symbol remains the same in the target, so we can therefore calculate the address of the symbol in the target process. Interesting symbols at this point are the function that returns whether the emulator is already running (*seal_isrunning()*), and the one starting the emulator (*seal_initandrun()*). Using their addresses we setup the SEAL activation shellcode before injecting it in the target.

At this point we read the target process's state that we need to pass to the emulator. It consists of the values of general purpose and floating point registers, as well as state used by MMX instructions. Finally, before proceeding to the next phase we inject the activation shellcode, in the ELF header of the executable which contains 240 bytes that remain unused after loading the binary into memory.

The second phase of possession starts by redirecting the target's execution flow to the beginning of the injected shellcode. The actions performed collectively by *Eudaemon* and the shellcode are shown in Figure 3.4, and can be summarised into the following: (1) check that the target is not already possessed, (2) allocate a memory block to be used as stack by the emulator library; (3) store the process state saved during the first phase in the memory block obtained in step 2; (4) call the initialisation and execution function of the emulator; (5) detach from the target process.

To avoid starting a possession procedure for an already possessed process, we first perform a call into the library to discover whether it is already running. The return value of the call is placed within the *eax* register. To retrieve the result, we place a trap instruction right after the call that returns control back to *Eudaemon*, where we can actually check whether we should proceed with the possession, or fallback reinstating the saved process state and detach.

Assuming that the process is not already possessed, execution resumes, and we attempt to allocate a memory area that will be used as a stack for the execution of the emulator. A new stack is necessary, since sharing the active stack between the emulator and the emulated code would lead to error. We use *mmap()* to request a new memory area from the OS, and verify its successful completion by using *ptrace()* semantics to receive control in *Eudaemon* right after the return of the system call.

Assuming control after the return of *mmap()* is also necessary to supply the required arguments to the emulator. The arguments comprise the process state that we read during the first phase of the possession, the exact state where native execution stopped. We inject the data into the newly

```

08048000-08049000 r-xp 00000000 03:04 4450978 loop
08049000-0804a000 rw-p 00000000 03:04 4450978 loop
40000000-40016000 r-xp 00000000 03:04 719528 /lib/ld-2.3.6.so
40016000-40018000 rw-p 00015000 03:04 719528 /lib/ld-2.3.6.so
40018000-40019000 r-xp 40018000 00:00 0 [vdso]
40019000-4001a000 rw-p 40019000 00:00 0
40034000-400c1000 r-xp 00000000 03:04 3140602 libseal.so.0.2
400c1000-400c9000 rw-p 0008c000 03:04 3140602 libseal.so.0.2
400c9000-42118000 rw-p 400c9000 00:00 0
42118000-4224a000 r-xp 00000000 03:04 719531 /lib/libc-2.3.6.so
42240000-42241000 r--p 00127000 03:04 719531 /lib/libc-2.3.6.so
42241000-42244000 rw-p 00128000 03:04 719531 /lib/libc-2.3.6.so
42244000-42246000 rw-p 42244000 00:00 0
42246000-42267000 r-xp 00000000 03:04 719535 /lib/libm-2.3.6.so
42267000-42269000 rw-p 00020000 03:04 719535 /lib/libm-2.3.6.so
bfa87000-bfa9d000 rw-p bfa87000 00:00 0 [stack]

```

Figure 3.5: Contents of a `/proc/[pid]/maps` file - note the presence of `libseal`

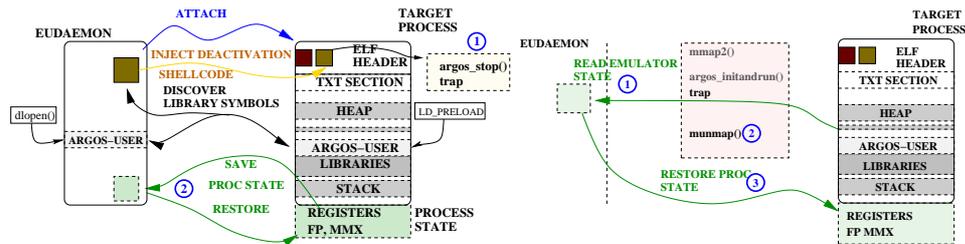


Figure 3.6: Process release: phase 1 Figure 3.7: Process release: phase 2

allocated stack, while also reducing its length by the size of the data being stored.

Placing the process state in the emulator stack is the last action performed by *Eudaemon*, which then detaches and exits. The shellcode within the target process performs the last step, and calls the emulator main routine, which initialises itself and starts the emulation.

3.3.2.2 Process Release

Releasing a process is also partitioned in two phases with the first being similar to possession. An overview is shown in Figure 3.6, and the *additional* steps in respect to possession (listed in Section 3.3.2.1) are the following: (1) call the emulator's stop routine, and at the same time discover whether it was running; (2) reinstate the saved process saved state, and allow it to resume execution.

Just like in possession, *Eudaemon* also attaches to the target process, looks up the required library symbols in the target, sets up the shellcode, and injects it. The additional assembly code introduced in the process does

not overlap with the shellcode injected during possession, and is quite small in size. It simply calls the *seal_stop()* function in the emulator, requesting it to exit. The same function also checks that the emulator is running, so there is no need to perform an additional call to retrieve its state beforehand.

If the process was possessed, *seal_stop()* initiates an exit from the emulator and reports success, while otherwise it returns error. We receive control back in *Eudaemon*, by inserting a trap instruction right after the call. We proceed to read its return value to determine whether the release request was valid, in which case *Eudaemon* waits for the emulator to exit. In any other case, it restores the saved process state allowing it to resume execution uninterrupted.

When the emulator exits, execution returns to the original shellcode planted during possession. The remainder of that code in conjunction with *Eudaemon* is responsible for switching a process's execution back to normal. Figure 3.7 shows an overview of this procedure, which in brief is: (1) recover the emulated process's state, stored in the emulator stack; (2) release the memory block that is used as stack; (3) restore the state read in step (1) as native process state.

As soon as the emulator exits, a trap instruction is executed to notify *Eudaemon* of the event. We then re-read the target's state to discover the address of the stack being used, and consequently the location of the emulator state that corresponds to the real process state we need to reinstate for release to be carried out. After recovering the state, the target is resumed and the stack we allocated is freed using *munmap()*. Once again, we use *ptrace()* semantics to receive control when this system call returns, to finally reinstate process state. Finally, we detach from the process effectively completing the release of the process.

3.4 Evaluation

We evaluate how *Eudaemon* performs in two aspects: the overhead induced on an application when executing under the emulator, and the cost of possessing and releasing.

3.4.1 SEAL

To evaluate the overhead imposed on an application when emulated by SEAL, we measured the performance of a set of UNIX programs when run natively and when SEAL-emulated. We also compare against the Argos full-system emulator. Our benchmark consists of one CPU-intensive application with little I/O (*bunzip2*), non-interactive network downloader with little CPU utilisation (*wget*), a network downloader with encryption (*sftp*), and one interactive graphical browser that performs both downloading and

	bunzip2	wget	sftp	konqueror
Native Execution	27.99s	10.97MB/s	14.3MB/s	29.4ms
SEAL (1 byte tags)	242.24s	10.92MB/s	2.3MB/s	463.4ms
Slowdown (<i>factor</i>)	×8.6	×1	×6.3	×15.6
Argos (1 byte tags)	508.66s	0.90MB/s	0.55MB/s	n/a
Slowdown (<i>factor</i>)	×18.2	×12.2	26	n/a
SEAL (1 bit tags)	248.78s	10.93MB/s	2.3	725ms
Slowdown (<i>factor</i>)	×8.9	×1	×6.3	×24.5
Argos (1 bit tags)	635.15s	0.49MB/s	0.47MB/s	n/a
Slowdown (<i>factor</i>)	×22.7	×22.4	26	n/a

Table 3.1: Emulation overhead

rendering (*konqueror*). Konqueror is the official web browser and file manager for KDE. With this mix of applications, we have covered the spectrum of use cases for *Eudaemon* fairly well so that the results represent a faithful indication of expected performance in general.

The experiments were conducted on a dual Intel™ Xeon at 2.80 GHz with 2 MB of L2 cache and 4 GB of RAM. The system was running SlackWare Linux 10.2 with kernel 2.6.15.4. The versions of the utilities used were *bzip2 v1.0.3*, *GNU wget v1.10.2*, and *konqueror 3.5.4*.

We used *bzip2* to decompress the Linux kernel 2.6.18 tar archive which amounts to about 40 MB of data. We used the UNIX utility *time* to measure the execution time of the decompression. For *wget*, and *sftp* we fetched the same file from a dedicated HTTP server over a 100 Mb/s LAN. In the experiment we used *wget* and *wget*'s own calculation of the average transfer rate as performance measure. Finally, we measured the time needed by *konqueror* to load and draw an HTML page along with a stylesheet. We used the `loadtime` browser benchmarking utility available from <http://nontroppo.org/test/Op7/loadtime.html> to conduct the measurement, but had it loaded locally to avoid incorporating variable network latencies in the experiment. Because of clock skew, a well-known problem with Qemu, we could not measure this test reliably on Argos. Table 3.1 shows the results.

We observe that compared to native execution *bunzip2* under SEAL requires about 8.5 times more time to complete. The overhead is fairly large, but this was expected and can be mainly attributed to the dynamic translator and the additional instrumentation. Nevertheless, it is much lower than the performance penalty suffered when using the Argos *system* emulator (i.e., if we run the entire OS on Argos), which compared to a native system was reported to run at least 16 to 20 times slower [29]. Furthermore, using *Eudaemon* we can choose *when* to employ emulation, reducing user inconvenience caused by the slowdown to a minimum.

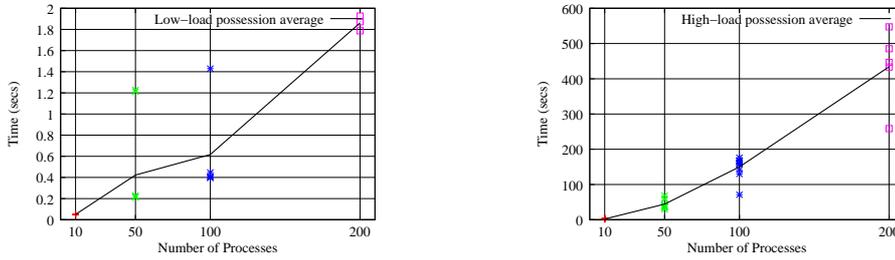


Figure 3.8: Multiple process possession under low- and high-load

Eudaemon Action	Possession	Release
1 st phase	1.195	0.159
Waiting time	<i>not applicable</i>	2782.617
2 nd phase	0.095	0.106
Total	1.290	2782.882

Table 3.2: Eudaemon micro-timings (msec)

The results from *wget* are quite different. The network transfer of a file was subject to insignificant performance loss. *Wget* performs no data processing, and the sole overhead is imposed by the instrumentation of *read* and *write* calls. The results are encouraging enough to allow for the possibility of running I/O dominated services such as FTP and file sharing entirely in emulation mode.

sftp incurs a slowdown of a factor 6.3. In our opinion, this is surprisingly good considering all the operations on tainted data involved in *ssh*. In other work, the reported overhead is more than two orders of magnitude [4]. We suspect that the difference is caused by the fact that *Eudaemon* attaches on the application after a shared secret key has already been established, and therefore does not suffer the initial expensive connection set up that uses asymmetric encryption.

Konqueror yields the worst results. We ascribe this to the fact that the GUI, as well as rendering the content, uses many instructions that incur much overhead in emulation, including floating point operations as well as MMX operations.

3.4.2 Eudaemon

Another important performance metric for *Eudaemon* is the time it takes to possess and consequently release a process. We examine these two operations from two various aspects. First we measure the time needed to possess and release a single process, by calculating the time spent on each of the two phases of the operations. Second we measure how process possession scales with an increasing number of targets.

```
while (honeypot_mode == true) {  
    run_instruction_block();  
    handle_system_call();  
    handle_signals();  
}
```

Figure 3.9: SEAL Execution Loop

Table 3.2 shows the total time needed for the possession and release of a single process, as well as how this time is distributed amongst the different phases as they were presented in Section 3.3. Possession of a single process takes very little time to complete. Release spends even less time performing the two phases, but it is delayed due to waiting for the emulator to exit gracefully. To clarify this point we present the main execution loop of SEAL in fig.3.9. After the completion of the second release stage, *Eudaemon* is blocked waiting for the current block of emulated instructions to conclude, and the emulator to exit its main loop. As a result the target process is not blocked during this time, and the observed delay is small.

To measure the performance of *Eudaemon* when multiple process are possessed, we created an increasing number of processes, which we proceed to possess. Figure 3.8 plots the time needed to switch a number of processes from native to emulated execution. The results also include the time needed to retrieve the PIDs of processes using *ps*, as well as to *fork()* a separate *Eudaemon* process to perform the possession for each target. The two graphs shown represent two different scenarios. In the left graph we possess idle processes that at the time of possession are within *sleep()*, while in the right graph we possess CPU intensive processes with 100% host CPU utilisation. Even though performance is lower in the latter, in both cases *Eudaemon* scales reasonably well. We believe that this experiment supports our claim that *Eudaemon*'s performance is suitable for the idle-time honeypots and honey-on-demand scenarios as presented in Section 3.

Regarding security, the emulator used in SEAL was tested against many types of exploit, including: Apache chunked encoding overflow, WebDav ntdll.dll overflow, IIS ISAPI .printer host header overflow, RPC DCOM Interface overflow, LSASS Overflow, nbSMTP remote format string exploit, NetApi exploit, WMF exploit, and many others. Interested readers are referred to the original Argos paper [29].

Chapter 4

Conclusions

We have described both *Shelia* and *Eudaemon*, two very different approaches towards detecting client-side attacks. *Shelia* is a 'traditional' client-side honeypot: a dedicated machine that can be used to scan lists of potentially malicious servers and documents. Compared to many existing honeypots it has an interesting detection method. Rather than scanning the file system for differences after an interaction, it flags a call to a sensitive operation as an attack if it was made from an area that should contain data. *Shelia* is fast and accurate, and it has proved effective against a host of real attacks. However, it cannot be used to protect production machines directly.

In contrast, *Eudaemon* is a technique that allows us to grab a running process and continue its execution in safe mode in an emulator. This way we can integrate the honeypot in the production machine. The *Eudaemon* emulator provides extensive instrumentation in the form of taint analysis to protect the application. It allows us to turn a machine into a honeypot in idle hours, or to protect applications that are about to perform actions that are potentially harmful. We have shown that the performance overhead of *Eudaemon* on Linux is reasonable for most practical use cases. To the best of our knowledge, this is the first system that allows one to force fully native applications to switch to emulation in mid-processing. We believe it provides an interesting instrument to increase the security of production machines. At the moment, *Eudaemon* is still a research prototype.

CHAPTER 4. CONCLUSIONS

Bibliography

- [1] Infecting elf-files using function padding for linux. <http://vx.netlux.org/lib/vhe00.html>.
- [2] Runtime process infection. <http://www.phrack.org/archives/59/p59-0x08.txt>.
- [3] Writing parasitic code in C. <http://ares.x25zine.org/ES/txt/C-parasites.txt>.
- [4] A. Ho, M. Fetterman, C. Clark, A. Warfield, and S. Hand. Practical taint-based protection using demand emulation. In *EuroSys, Leuven, Belgium*, April 2006.
- [5] K. G. Anagnostakis, S. Sidiroglou, P. Akritidis, E. M. K. Xinidis, and A. D. Keromytis. Detecting targeted attacks using shadow honeypots. In *Usenix Security*, Baltimore, MD, August 2005.
- [6] F. Bellard. QEMU, a fast and portable dynamic translator. In *Proc. of USENIX ATC*, pages 41–46, Anaheim, CA, April 2005.
- [7] S. Bhatkar, D. D. Varney, and R. Sekar. Address obfuscation: an efficient approach to combat a broad range of memory error exploits. In *Proc. of USENIX Security*, pages 105–120, August 2003.
- [8] D. Brumley, J. Newsome, D. Song, H. Wang, and S. Jha. Towards automatic generation of vulnerability-based signatures. In *Security and Privacy*, Oakland, CA, May 2006.
- [9] C. Cowan, S. Beattie, J. Johansen and P. Wagle. PointGuard: Protecting pointers from buffer overflow vulnerabilities. In *In Proc. of the 12th USENIX Security Symposium*, pages 91–104, August 2003.
- [10] J. R. Crandall and F. T. Chong. Minos: Control data attack prevention orthogonal to memory model. In *Proc. of the 37th annual International Symposium on Microarchitecture*, pages 221–232, Portland, Oregon, 2004.
- [11] Crispin Cowan, Calton Pu, Dave Maier, Heather Hintony, Jonathan Walpole, Peat Bakke, Steve Beattie, Aaron Grier, PerryWagle and Qian Zhang. StackGuard: Automatic Adaptive Detection and Prevention of Buffer-Overflow Attacks. In *7th USENIX Security Symposium*, San Francisco, CA, 2002.
- [12] W. Cui, V. Paxson, N. Weaver, and R. Katz. Protocol-independent adaptive replay of application dialog. In *The 13th Annual Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, February 2006.
- [13] W. de Bruijn, A. Slowinska, K. van Reeuwijk, T. Hruby, L. Xu, and H. Bos. Safecard: a gigabit ips on the network card. In *Proceedings of 9th International Symposium on Recent Advances in Intrusion Detection (RAID'06)*, pages 311–330, Hamburg, Germany, September 2006.
- [14] D. Denning. A lattice model of secure information flow. *ACM Transactions on Communications*, 19(5):236–243, 1976.

BIBLIOGRAPHY

- [15] eEye. eeye industry newsletter. <http://www.eeye.com/html/resources/newsletters/versa/VE20070516.html#techtalk>, May 2007.
- [16] W. W. Hsu and A. J. Smith. Characteristics of i/o traffic in personal computer and server workloads. *IBM Systems Journal*, 42(2), 2003.
- [17] G. S. Kc, A. D. Keromytis, and V. Prevelakis. Countering code-injection attacks with instruction-set randomization. In *Proc. of the ACM Computer and Communications Security (CCS)*, pages 272–280, Washington, DC, October 2003.
- [18] N. Krawetz. Anti-honeypot technology. *IEEE Security and Privacy*, 2(1):76–79, January 2004.
- [19] B. Lampson. Accountability and freedom. In *Cambridge Computer Seminar*, Cambridge, UK, October 2005.
- [20] C. Leita, M. Dacier, and F. Massicotte. Automatic handling of protocol dependencies and reaction to 0-day attacks with scriptgen based honeypots. In *Proceedings of RAID'06*, pages 185–205, Hamburg, Germany, September 2006.
- [21] M. E. Locasto, S. Sidiroglou, and A. D. Keromytis. Application communities: Using monoculture for dependability. In *Proceedings of the 1st Workshop on Hot Topics in System Dependability (HotDep)*, pages 288 – 292, Yokohama, Japan, June 2005.
- [22] M. E. Locasto, A. Stavrou, G. F. Cretu, and A. D. Keromytis. From stem to sead: Speculative execution for automated defense. In *Proceedings of the 2007 USENIX Annual Technical Conference*, pages 219–232.
- [23] M. Costa, J. Crowcroft, M. Castro, A. Rowstron, L. Zhou, L. Zhang and P. Barham. Vigilante: End-to-end containment of internet worms. In *In Proc. of the 20th ACM Symposium on Operating Systems Principles (SOSP)*, Brighton, UK, October 2005.
- [24] A. Moshchuk, T. Bragin, S. D. Gribble, and H. M. Levy. A crawler-based study of spyware in the web. In *Proc. of NDSS'06*, San Diego, CA, February 2006.
- [25] J. Newsome, D. Brumley, J. Franklin, and D. Song. Replayer: automatic protocol replay by binary analysis. In *CCS '06: Proceedings of the 13th ACM conference on Computer and communications security*, pages 311–321, New York, NY, USA, 2006. ACM Press.
- [26] J. Newsome and D. Song. Dynamic taint analysis for automatic detection analysis and signature generation of exploits on commodity software. In *Proc. of the 12th Annual Network and Distributed System Security Symposium (NDSS)*, 2005.
- [27] A. One. Smashing the stack for fun and profit. *Phrack Magazine*, 7(49), 1996.
- [28] G. Portokalidis and H. Bos. Eudaemon: Involuntary and on-demand emulation against zero-day exploits. In *Proceedings of ACM SIGOPS EUROSYS'08*, pages 287–299, Glasgow, Scotland, UK, April 2008. ACM SIGOPS.
- [29] G. Portokalidis, A. Slowinska, and H. Bos. Argos: an Emulator for Fingerprinting Zero-Day Attacks. In *Proc. of the 1st ACM SIGOPS EUROSYS*, Leuven. Belgium, April 2006.
- [30] J. Richter. Load your 32-bit dll into another process's address space using injlib. *Microsoft Systems Journal (MSJ)*, January 1996.
- [31] SANS. Sans institute press update. http://www.sans.org/top20/2006/press_release.pdf, 2006.
- [32] H. Shacham, M. Page, B. Pfaff, E. Goh, and N. Modadugu. On the effectiveness of address-space randomization. In *Proceedings of the 11th ACM conference on Computer and communications security*, pages 298–307. ACM, 2004.
- [33] S. Sidiroglou, M. Locasto, S. Boyd, and A. Keromytis. Building a reactive immune system for software services. In *Proceedings of the 2005 USENIX Annual Technical Conference*, 2005.

- [34] A. Slowinska and H. Bos. The age of data: pinpointing guilty bytes in polymorphic buffer overflows on heap or stack. In *23rd Annual Computer Security Applications Conference (ACSAC'07)*, Miami, FLA, December 2007.
- [35] P. Szor and P. Ferrie. Hunting for metamorphic. In *Virus Bulletin Conference*, pages 123–144, Abingdon, Oxfordshire, England, September 2001.
- [36] J. Tucek, S. Lu, C. Luang, S. Xanthos, Y. Zhou, J. Newsome, D. Brunmley, and D. Song. Sweeper: a light-weight end-to-end system for defending against fast worms. In *Proceedings of Eurosys 2007*, Lisbon, Portugal, April 2007.
- [37] X. Wang, Z. Li, J. Xu, M. K. Reiter, C. Kil, and J. Y. Choi. Packet vaccine: black-box exploit detection and signature generation. In *CCS '06: Proceedings of the 13th ACM conference on Computer and communications security*, pages 37–46, New York, NY, USA, 2006. ACM Press.
- [38] Y.-M. Wang, D. Beck, X. Jiang, R. Roussev, C. Verbowski, S. Chen, and S. King. Automated web patrol with strider honeymonkeys: Finding web sites that exploit browser vulnerabilities. In *Proc. Network and Distributed System Security (NDSS)*, San Diego, CA, February 2006.
- [39] Wikipedia. Entry on call stacks, 2007.
- [40] S. M. B. William R. Cheswick, Aviel D. Rubin. *Firewalls and Internet Security: repelling the wily hacker (2nd ed.)*. Addison-Wesley, ISBN 020163466X, 2003.
- [41] C. C. Zou and R. Cunningham. Honeypot-aware advanced botnet construction and maintenance. In *The International Conference on Dependable Systems and Networks (DSN-2006)*, Philadelphia, PA, USA, June 2006.